

Petabyte Databases

Dirk Düllmann
CERN IT/ASD & RD45
Geneva

+41 22 767 4937

Dirk.Duellmann@cern.ch

ABSTRACT

This paper describes the use of Object-Database Management Systems (ODBMS) for the storage of High-Energy Physics (HEP) data.

Keywords

Object-Databases, Very Large Databases, High Energy Physics

1. Physics Data Management

The new experiments at the Large Hadron Collider (LHC) at CERN will gather an unprecedented amount of data. Starting from 2005 each of the four LHC experiments ALICE, ATLAS, CMS and LHCb will measure of the order of 1 Peta Byte (10^{15} Bytes) per year. All together the experiments will store and repeatedly analyse some 100 PB of data during their lifetimes. Such an enormous task can only be accomplished by large international collaborations. Thousands of physicists from hundreds of institutes world-wide will participate. This also implies that nearly any available hardware platform will be used resulting in a truly heterogeneous and distributed system.

The computing technical proposals of LHC experiments do not only require access to the data from remote sites but in addition ask for distribution of the data store itself to several regional centers.

Experiment	Data Rate	Data Volume
ALICE	1.5 GB/sec	1 PB/month (1 month per year)
ATLAS	100 MB/sec	1 PB/year
CMS	100 MB/sec	1 PB/year
LHCb		400 TB/year

Table 1: Expected Data Rates and Volumes at LHC

2. Physics Data Models

Data Models in High Energy Physics (HEP) are typically very complex. The number of different data types needed to describe the event data of a large HEP experiment reaches several hundreds. A single measurement (event) contains up to millions of interrelated objects.

Since the object model of is shared between multiple subsystems (e.g. data acquisition, event reconstruction and physics analysis) with very different access patterns, it is often difficult to fulfil all flexibility and performance requirements in a single design.

All LHC experiments exploit Object Oriented (OO) technology to implement and maintain their very large software systems. Today most software development is done in C++ with a growing interest in Java. The data store therefore has to support the main concepts of these OO languages such as abstraction, inheritance, polymorphism and parameterised types.

3. Conclusion

HEP data stores based on Object Database Management Systems (ODBMS) are very well suited to support HEP object models. They provide a number of important advantages in comparison to traditional systems like a single coherent logical view of complex HEP object models and allows a tight integration with multiple of today's OO languages.

The clear separation of logical and physical data model allows for transparent support of physical clustering and re-clustering which is expected to be an important tool to optimise the overall system performance.

The ODBMS implementation of Objectivity/DB in particular show scaling up to multi-PB distributed data stores and provides a seamless integration with Mass Storage Systems (MSS) like HPSS. Already today a significant number of HEP experiments in or close to production adopted an ODBMS based approach

4. REFERENCES

- [1] RD45 - A Persistent Object Manager for HEP, LCB Status Report, March 1998, CERN/LHCC 98-
- [2] RD45 - A Persistent Object Manager for HEP, LCB Status Report, March 1997, CERN/LHCC 97-6
- [3] Using and Object Database and Mass Storage System for Physics Analysis, the RD45 collaboration
- [4] RD45 - A Persistent Object Manager for HEP, LCRB Status Report, March 1996, CERN/LHCC 96-15

[5] Object Databases and Mass Storage Systems: The Prognosis, the RD45 collaboration, CERN/LHCC 96-17

and references therein